arXiv:1404.1890v1 [cs.CL] 7 Apr 2014

Polish and English wordnets - statistical analysis of interconnected networks

Maksymilian Bujok, Piotr Fronczak, Agata Fronczak

Faculty of Physics, Warsaw University of Technology, Koszykowa 75, PL-00-662 Warsaw, Poland

Wordnets are semantic networks containing nouns, verbs, adjectives, and adverbs organized according to linguistic principles, by means of semantic relations. In this work, we adopt a complex network perspective to perform a comparative analysis of the English and Polish wordnets. We determine their similarities and show that the networks exhibit some of the typical characteristics observed in other real-world networks. We analyse interlingual relations between both wordnets and deliberate over the problem of mapping the Polish lexicon onto the English one.

PACS numbers: 89.65.Gh,89.75.-k,05.40.-a,02.10.Ox

1. Introduction

A wordnet is a network of concepts which are connected according to their meaning by semantic relations. Its importance for natural language processing is already acknowledged through hundreds of projects and tools that relay on it [1]. The most important wordnet applications include word sense disambiguation [2], language teaching and translation [3], information retrieval [4], and document classification [5].

The first, and by far the best developed lexical database of this kind is the one created for the English language and developed at the University of Princeton - the so called Princeton WordNet (PWN) [6]. It is commonly used as a reference for other lexical networks (several dozens of wordnets that follow PWN design are listed by the Global WordNet Associaton [7]).

Currently, huge efforts are being undertaken to map PWN onto other languages in aim to create a global wordnet grid as a one multilingual system [8]. For this reason, the understanding of the structure of particular wordnets, their differences and similarities, is not only interesting on its own right, as a challenging scientific problem. It is important because it

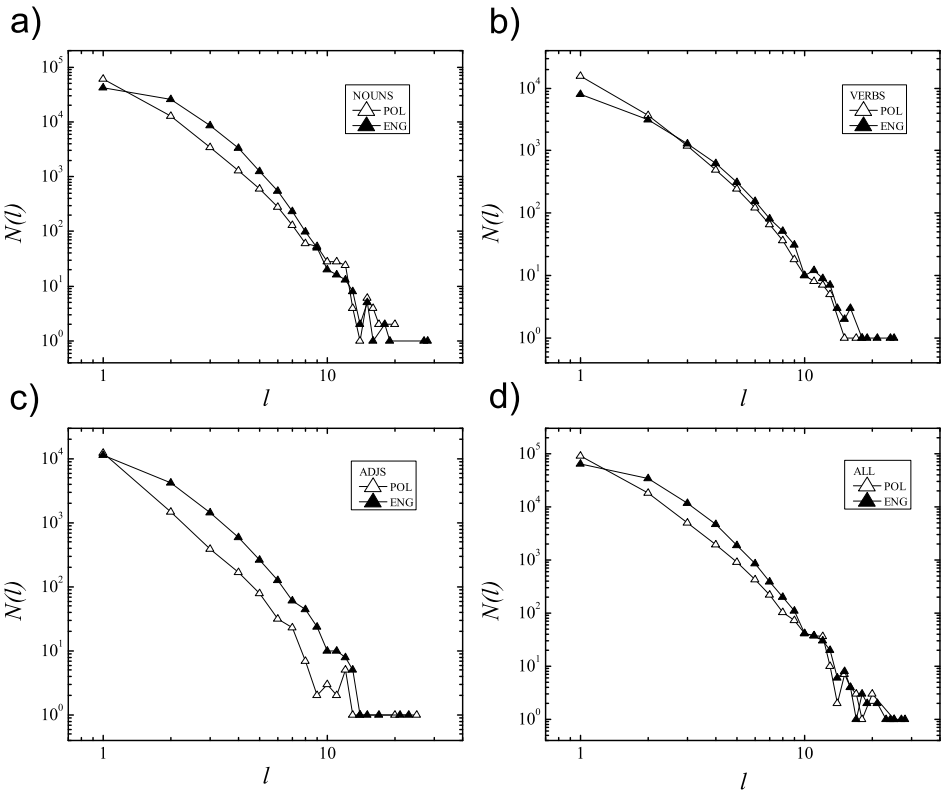

Fig. 1. The number of synsets composed of l lexems in both wordnets; a) nouns, b) verbs, c) adjectives, and d) all the grammatical categories.

provides an insight into the traps that researchers can fall into while semi-automatically matching objects from wordnets in unrelated languages.

In this paper we study several statistical characteristics of PWN and Polish wordnet (PolWN) [9] using tools from the field of complex networks. We compare and discuss similarities and differences between both systems. Finally, we propose a method to assess the quality of mapping of both wordnets taking into account existing inter-lingual relations between them. We point out that due to differences in local topology of different wordnets the local network measures (e.g. node degree, clustering coefficient, motifs, etc.), which are frequently used in studying lexical networks [10, 11, 12, 13], can be misleading.

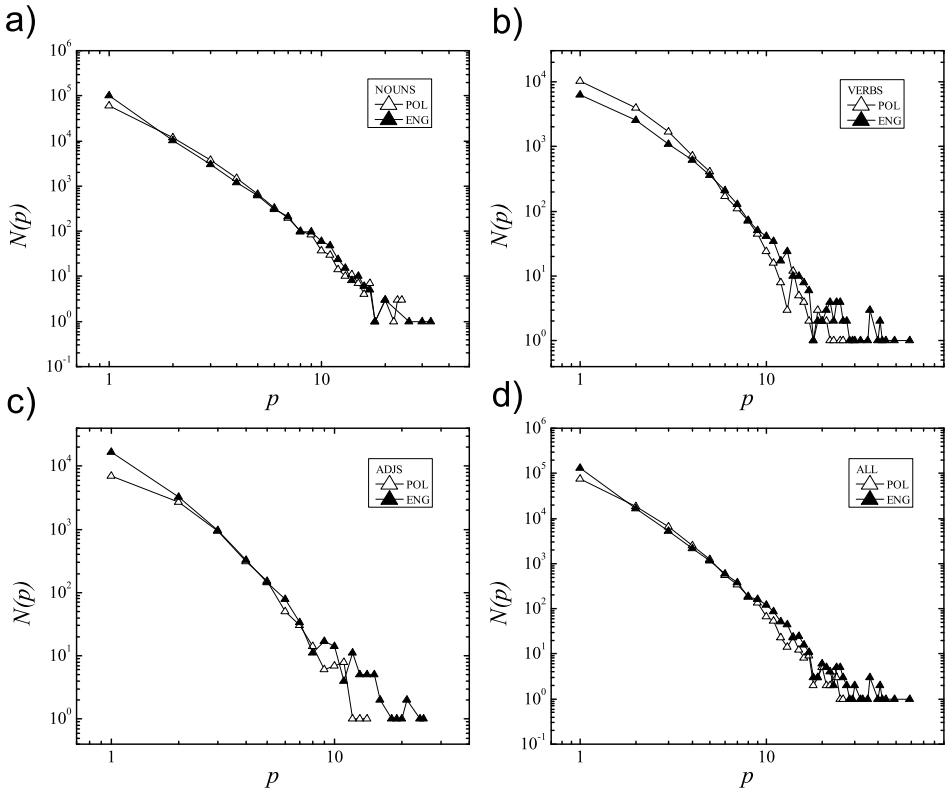

Fig. 2. The polysemy distribution of the Polish and English lexems (the number of lexems that have the given number of senses); a) nouns, b) verbs, c) adjectives, and d) all the grammatical categories.

2. Basic dataset statistics

The basic building block of a wordnet is a *synset*. Synsets are sets of synonyms that gather lexical items (*lexems*) being interchangeable in a given context without changing the meaning. In the analyzed PWN and PolWN databases, there are 117659 and 116319 synsets, and 206978 and 160098 lexical items, respectively. It means that the average number of lexems in a synset is quite small (less than two). However, one should note, that in both wordnets, and for different grammatical forms (nouns, verbs, and adjectives), the synset size distributions are fat-tailed (see Fig. 1). Therefore, although more than half of all synsets are composed of only one lexem, the largest synsets contain 25 (PolWN) and 28 (PWN) synonymous lexems. Despite apparent similarities between both lexicons, Fig. 1a and 1c may suggest that the synonymy in PolWN is more restrictive than in PWN

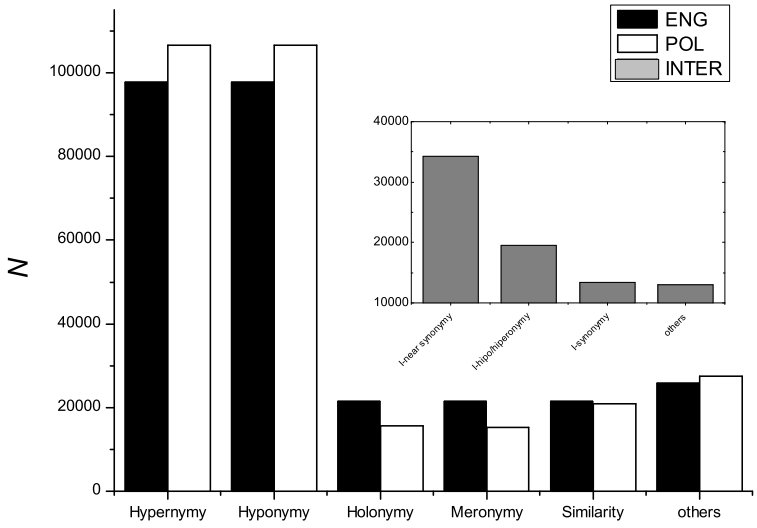

Fig. 3. Instances of semantic relations in the Polish and English wordnets. Inset shows the most important inter-lingual relations.

[8] or that the Polish wordnet is still relatively less mature than PWN.

Similarly, as a synset may contain many lexems (synonymy), a lexeme can participate in many different synsets (polysemy). Polysemy, the phenomenon when a single word has two or more meanings, is strictly related with the economy of the language, metaphoric thinking, imagery, and generalization [14]. On the other hand, the existence of polysemous words remains a central challenge to artificial intelligence in word sense disambiguation algorithms [15]. Fig. 2 presents the polysemy distributions of the Polish and English lexems with respect to different grammatical categories. In all the cases, there are no significant differences between both wordnets. However in the range of extreme values, English polysemous lexems are more frequent. With respect to this remark, if one considers a wordnet as a network where lexems are connected when they belong to the same synset, then such extremely polysemous words act as hubs. It is well know, that highly connected nodes shorten distances within networks [16]. Now, since distance in wordnets can be a useful proxy for semantic relatedness [22], one can conclude that word sense disambiguation in the English lexicon, having more hubs, is more prone to errors. This hypothesis should be verified in a future.

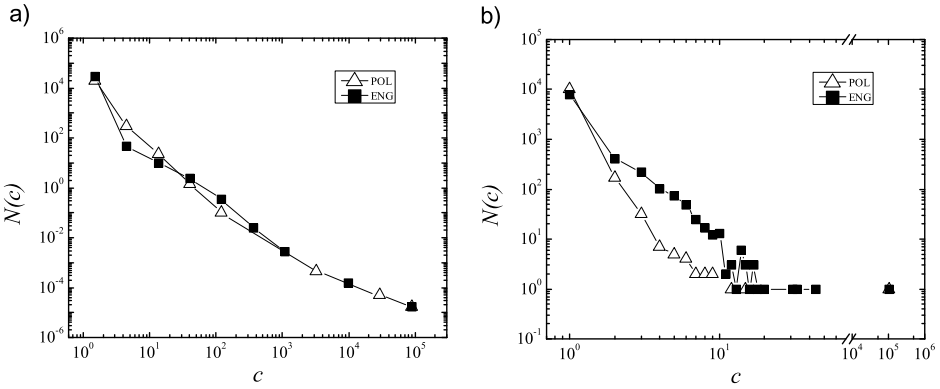

Fig. 4. Cluster size distributions for the Polish and English networks with a) hyperonymy links; b) all types of relations. Points in a) are averages over logarithmic bins along the horizontal axis.

3. Wordnet as a complex network

Synsets can be connected to other synsets via a number of semantic relations. The most important relationship is the hypernym/hyponym relation which bind a given concept to a more general/specific concept creating hierarchical connection, e.g. a *dog* is a hyponym of *animal* and *animal* is a hyperonym of *dog*. Another hierarchical relationship worth to mention is the meronym/holonym relation, which denotes part-whole/whole-part relationships, e.g. an *engine* is a part of a *car* (meronymy), and a *car* contains an *engine* (holonymy). There are several other relations (e.g. antonymy, casuality, etc.), but the semantic network is strongly dominated by the hypernymy/hyponymy links (see Fig. 3). That is why, in further analysis, we will take into account only this type of relationship.

The synsets connected by hyperonymy links form an acyclic directed graph composed of a large number of isolated clusters. It is worth to mention that these clusters are not trees as was stated in the influential paper by Sigman and Cechi [14]. If one neglects directed character of the links, loops become quite frequent. The largest clusters in both wordnets (PolWN and PWN) consist of the nouns and contain respectively 59% and 64% of all synsets. The fat-tailed character of the cluster size distribution is shown in Fig. 4a. In this context both wordnets are very similar. To show the differences, one has to take into account all available relations. In that case, one can expect that the separated clusters merge together into the one giant cluster, i.e. there are no concepts isolated from the rest of the semantic network. Surprisingly, it is not true in both languages. Fig. 4b presents the

cluster size distributions for the networks constructed as described above (with all relations). There exists the largest cluster (of size of 105661 and 106236 in PolWN and PWN respectively), but there are also many small-size clusters separated form the main core of the language. The most reliable explanation of this fact is that the construction of both wordnets is still not completed. If this explanation is correct, then, surprisingly, the younger Polish database outperforms the English one in that context.

Importance of a synset can be estimated on several criteria, including the number of relations with other synsets, a position in the hyperonymic hierarchy of concepts [18], or a sum of the frequencies of its component lexical units in the language corpus [17]. Since we have no access to the Polish corpus database, we decided to study synset importance with the help of the supremacy - a measure which encapsulates the first two criteria. Supremacy was proposed in [19] as a parameter that can play an important role for description of a class of directed networks. It describes the number of nodes that are subordinated to a certain node. The parameter equals to the size of the basin connected to a certain node or to the size of its in-component and as such it has been used for the description of the Internet structure [20] and for the scaling relations in food webs [21]. An example of the network with the calculated supremacies of its nodes is presented in Fig. 5. With respect to this figure, please note, that the node with the supremacy equal to 5 has less direct relations with other synsets than the node with the supremacy equal to 3. Moreover, both nodes have the same position in the hierarchy. It demonstrates outperformance of the supremacy measure over standard criteria. Fig. 6 presents the supremacy distribution for both wordnets. As in other complex networks, it is given by a power law. In both cases the characteristic exponent equals 1.8. In the next section, we will take advantage of the similarity of both distributions when analyzing the quality of mapping PolWN onto PWN.

In Fig. 7 the relation between the average supremacy and the synset size (i.e. the number of the lexems the synset is composed of) is presented. Since the supremacy distribution is fat-tailed, instead of classical arithmetic averaging over the synsets with the same size, geometric averaging has been used. For the major parts of the lexicons (i.e. for more than 99% of synsets in each wordnet) there are clear linear relations between the logarithm of the supremacy and the synset size, i.e. $\langle s \rangle(l) \propto \exp(\alpha l)$. Interestingly, the scaling exponents α differ significantly between both wordnets ($\alpha_{POL} = 0.26$ and $\alpha_{ENG} = 0.08$). With respect to this observation, one can hypothesize that a new synonym, which enters into the synset when the language evolves, makes that concept much more important in Polish language than in English.

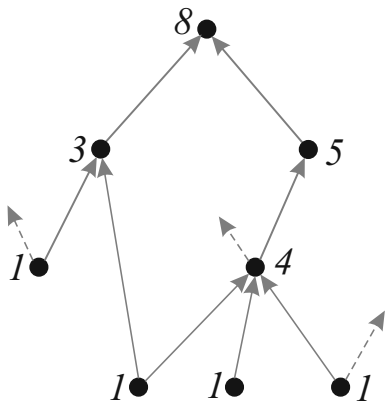

Fig. 5. Schematic illustration of the supremacy measure. Solid arrows represent connections within the supremacy area of the top vertex, whereas dashed arrows express connections pointing outside that area. Numbers situated in the vicinity of the nodes represent their supremacies.

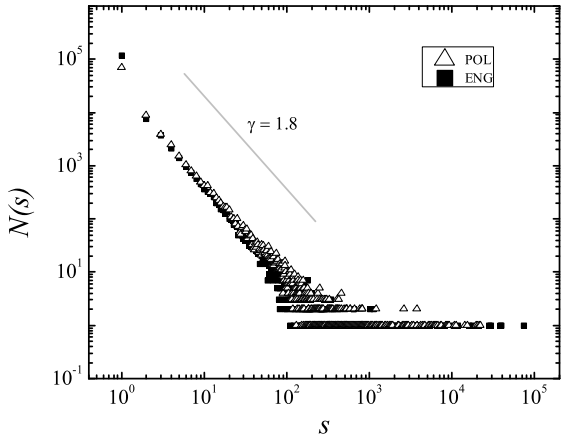

Fig. 6. Supremacy distributions in PWN and PolWN. Solid line represents the power law slope of both distributions.

4. The analysis of inter-lingual relations

Interconnected network systems are of interest to researchers in numerous different disciplines [23, 24, 25]. In linguistics the classical example of such a system is the EuroWordNet - the database that stores several European wordnets as autonomous language-specific structures that are inter-

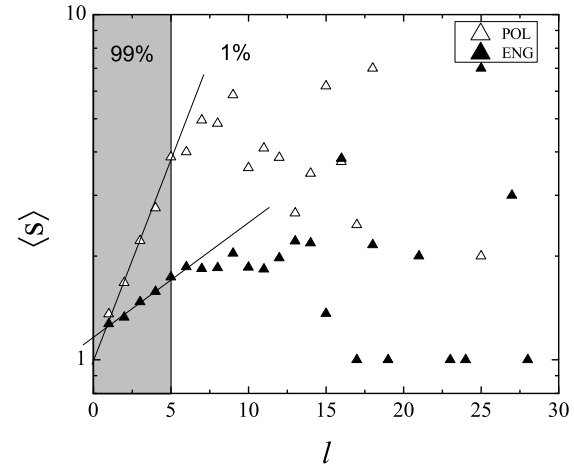

Fig. 7. The relation between the average supremacy of the synset and its size (i. e. the number of the lexems the synset is composed of). The straight lines are used to emphasize the linear character of the relation. Gray region contains more than 99% of the synsets in both wordnets.

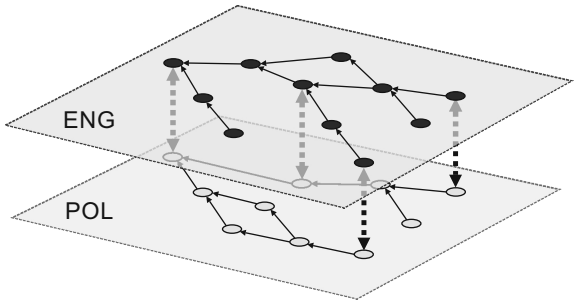

Fig. 8. The two-layered network composed of the Polish and English wordnets with the hiperonymy (thin arrows) and i-synonymy (thick dotted arrows) relations.

connected via inter-lingual relations [26]. Since a wordnet-wordnet mapping holds great potential for crosslinguistic applications, the statistical analysis of the interconnected wordnets can shed a light on the commonalities and differences in the ways languages map concepts onto words. For example, it can help to detect language-specific lexical gaps, where a word in one language has no correspondence in another language [27].

Problems with the mapping of PolWN onto PWN arise from both the differences in the conceptual and lexico-grammatical structure of English

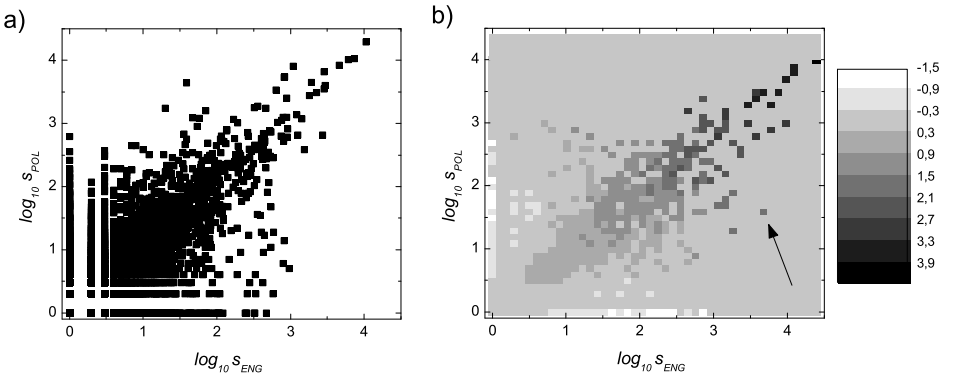

Fig. 9. Supremacy profiles of the sysnets related by the inter-lingual synonymy. The colors represent the value of the ratio $R(s_{ENG}, s_{POL})$, which detailed description can be found in the text.

and Polish languages, as well as from different methodological assumptions which underlie the construction of plWordNet where, on the contrary to PWN, lexems, instead of synsets, are the basic building elements [8].

Fig. 8 presents the two-layered network composed of the Polish and English wordnets (directed acyclic graphs) connected by the inter-lingual synonymy (the so-called i-synonymy) relations. It is worth noting that, in general, the set of relations that connect both wordnets includes many different types of relations, such as inter-lingual synonymy and near-synonymy, inter-lingual hyponymy and hypernymy, and inter-lingual meronymy and holonymy, and some of them are much more frequent than i-synonymy (see Fig. 3). However, i-synonymy represents the most direct correspondence between the senses of the connected synsets. Therefore, for the sake of clarity we only take into account i-synonymy relations since they seem to be the most prone to matching errors.

The analyzed network contains 117659 and 116319 nodes, 89087 and 106549 directed links in the English and Polish layers respectively. Moreover, $L = 13336$ i-synonymy links connect both layers. Now, it is reasonable to expect that the matched synsets share similar semantic importance in their own wordnets. Thus, when comparing their supremacies one should observe a clear correlation pattern. Each deviation from such a pattern can signal a potential mismatch of synsets or tell us about a language-specific gap in one of the two lexicons. Fig. 9a confirms that such a pattern exists and especially in the range of the large supremacies the correlation is clearly visible. On the other hand, in Fig. 9a, in the range of small suprema-

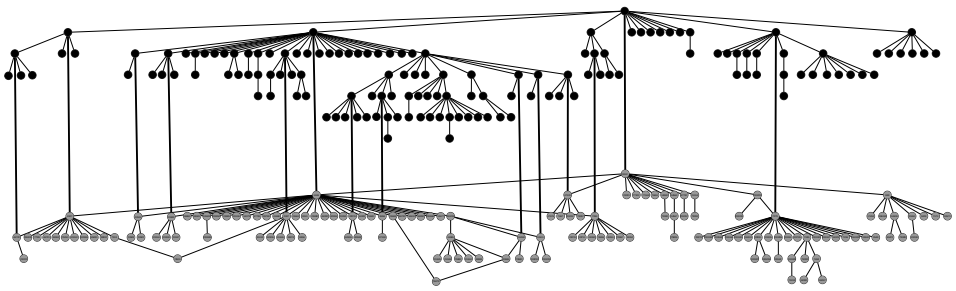

Fig. 10. The in-components of the Polish synset *publikacja* (gray nodes) and the English sysnet *publication* (black nodes). Vertical lines represent inter-lingual synonymy relations.

cies a dispersion of points is high and one may get the impression that the topology of the i-synonymy connections is similar to that created in a completely random way. Such a thinking is however not correctly enough. To see this, let $L(s_{ENG}, s_{POL})$ denotes the total number of links connecting synsets with supremacies s_{ENG} and s_{POL}. This number can be compared to its typical value $L_0(s_{ENG}, s_{POL})$ in the appropriate null-model network, i.e. the network with the randomly distributed i-synonymy links [28]. In such a network

$$L_0(s_{ENG}, s_{POL}) = p(s_{ENG})p(s_{POL})L, \tag{1}$$

where $p(s)$ is the probability that a synset has a supremacy s (the both probabilities, $p(s_{ENG})$ and $p(s_{POL})$, can be calculated, e.g., from Fig. 6). In Fig. 9b we illustrate the above comparison by plotting the ratio

$$R(s_{ENG}, s_{POL}) = \log_{10} \frac{L(s_{ENG}, s_{POL})}{L_0(s_{ENG}, s_{POL})}. \tag{2}$$

In our analysis, supremacies characterizing PWN and PolWN synsets are logarithmically binned into five bins per decade. Regions, where $R(s_{ENG}, s_{POL})$ is above (below) 0 correspond to enhanced (suppressed) connections between English and Polish synsets as compared to the randomized network. The presented procedure clearly demonstrates a linear pattern of correlations between the matched synsets. A better agreement for the larger values of the supremacy suggests that the matching has higher accuracy for the more general synsets. For example, Fig. 10 presents two in-components of the Polish synset {*publikacja, wydawnictwo*} and the English synset {*publication*}. The supremacies of both synsets are 141 and 140 respectively. Thick vertical lines represent existing inter-lingual synonymy relations between the

corresponding synsets. Although there is a visible similarity between both in-components, we would like to point out that they have different local topologies (which are probably due to different construction procedures of both wordnets). Thus, the two structures should not be compared on the level of node degrees, clustering coefficients, or other local measures as it was done in the previous studies (cf. [10, 11, 12, 13]). The supremacy is an example of a measure which also takes into account more distant neighborhood of the node and as such it outperforms simple local measures. Finally, an example of the potentially incorrect mapping is marked by the arrow in Fig. 9. The Polish synset {*wiedza*} has a supremacy 39 while its matched English counterpart {*cognition, knowledge*} - 4380.

5. Conclusions

We have analyzed several statistical properties of the Polish and English wordnets with the help of the tools from the field of complex networks. We have proposed the supremacy as a simple but not yet trivial measure of the importance of synsets. Finally, we have used it for the analysis of the matching of both lexicons. It has been suggested that the presented methodology can be used to assess the matching quality or to find language-specific lexical gaps.

6. Acknowledgements

This work was supported by the Foundation for Polish Science (grant no. POMOST/2012-5/5).

REFERENCES

[1] http://lit.csci.unt.edu/∼ wordnet/

[2] T. C. Silva et al., *Europhys. Lett.* **98**, 58001 (2012)

[3] X. Hu, A. Graesser, *Using WordNet and latent semantic analysis to evaluate the conversational contributions of learners in tutorial dialogue*, Proceedings of ICCE'98, 2. Beijing, China Higher Education Press, 337341 (1998).

[4] J. Y. Nie, M. Brisebois, *Articial Intelligence Review*, **10**, 409439 (1996).

[5] X. Peng, B. Choi, *Document classifications based on word semantic hierarchies*, Proceedings of the International Conference on Artificial Intelligence and Applications (AIA'05), 362–367 (2005).

[6] G. Miller, *Int. J. Lexicogr.* **3**, 245–264 (1990).

[7] http://globalwordnet.org/

[8] E. Rudnicka, et. al., *A strategy of mapping Polish WordNet onto Princeton Wordnet*, Proceedings of COLING 2012, Mumbai, India, 1039–1048 (2012).

[9] M. Piasecki, S. Szpakowicz, B. Broda, *A Wordnet from the Ground Up*, (Oficyna Wydawnicza Politechniki Wrocławskiej, Wrocław 2009).

[10] J. Li, et. al., *Physica A* **391**, 5254–5263 (2012).

[11] M. Steyvers, J. B. Tenenbaum, *Cognitive Sci.* **29**, 41–78 (2005).

[12] J. Ramanand, et. al., *IEEE Data Engineering Bulletin* **30**, 30–43 (2007).

[13] J. Borge-Holthoefer, A. Arenas, *Entropy* **12**, 1264–1302 (2010).

[14] M. Sigman, G. A. Cecchi, *Proc. Nat. Acad. Sci.* **99**, 1742–1747 (2002).

[15] Y. Ravin, C. Leacock, *Polysemy: Theoretical and Computational Approaches*, (Oxford Univ. Press, Oxford, 2000).

[16] J. A. Hołyst, et. al., *Physica A* **351**, 164 (2005).

[17] P. Buitelaar, B. Sacaleanu, *Extending Synsets with Medical Terms*, Proceedings of the First International Conference on Global Wordnet, Mysore, India, 21–25 (2002).

[18] P. Vossen, et. al., *The EuroWordNet Base Concepts and Top Ontology*, http://www.ncstrl.org:8900/ncstrl/servlet/search?formname=detail&id=oai%3Ancstrlh%3Aercimcnrcnuce%3Aercim.cnr.ilc%2F%2F1998-TR-004 (Centre National de la Recherche Scientifique, Paris, France, 1998).

[19] J. A. Hołyst, A. Fronczak, P. Fronczak, *Phys. Rev. E* **70**, 046119 (2004).

[20] G. Caldarelli, R. Marchetti and L.Pietronero, *Europhys. Lett.* **52**, 386 (2000).

[21] D. Garlaschelli, G. Caldarelli and L. Pietronero, *Nature* **423**, 165 (2003).

[22] A. Budanitsky, G. Hirst, *Comput. Linguist.* **32**, 13–47 (2006).

[23] S. N. Dorogovtsev, et. al., *Phys. Rev. E* **78**, 056106 (2008).

[24] S. V. Buldyrev, et. al., *Nature* **464**, 08932 (2010).

[25] J. Gao, et. al., *Phys. Rev. Lett.* **107**, 195701 (2011).

[26] P. Vossen, *Int. J. Lexicogr.* **17**, 161–173 (2004).

[27] A. Pease, C. Fellbaum, *Formal ontology as interlingua: the SUMO and WordNet linking project and global WordNet linking project and global WordNet.* In: *Ontology and the Lexicon: A Natural Language Processing Perspective*, ch.2., (Cambridge University Press, 2010).

[28] S. Maslov, K. Sneppen, A. Zaliznyak, *Physica A* **333**, 529–540 (2004).